



TAC 2010 Workshop

Slot Filling through Statistical Processing and Inference Rules

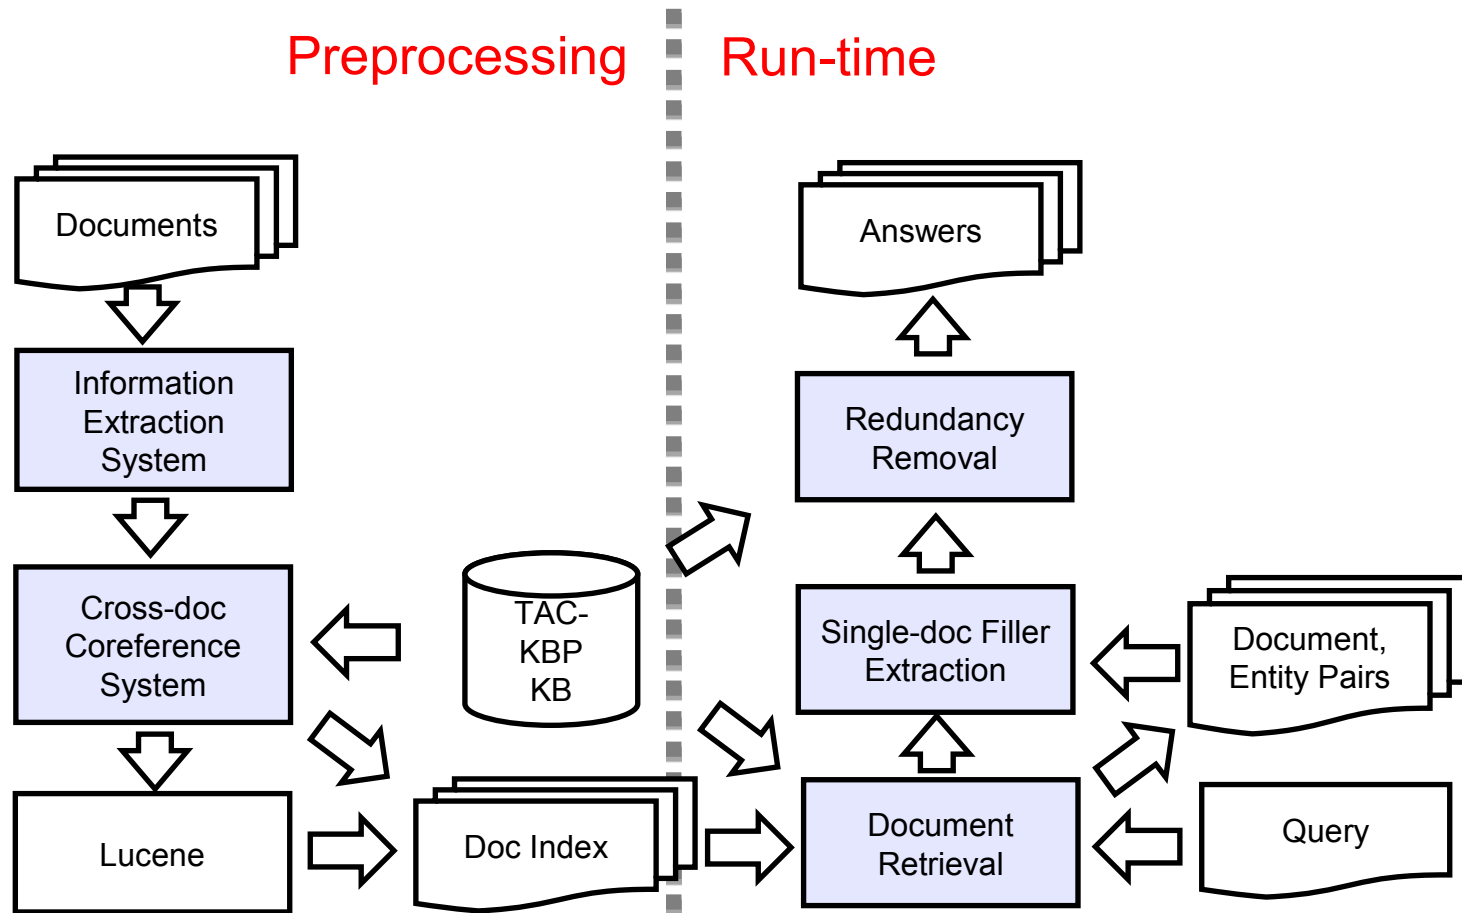
Vittorio Castelli, Radu Florian, Ding-Jung Han

November 15th, 2010

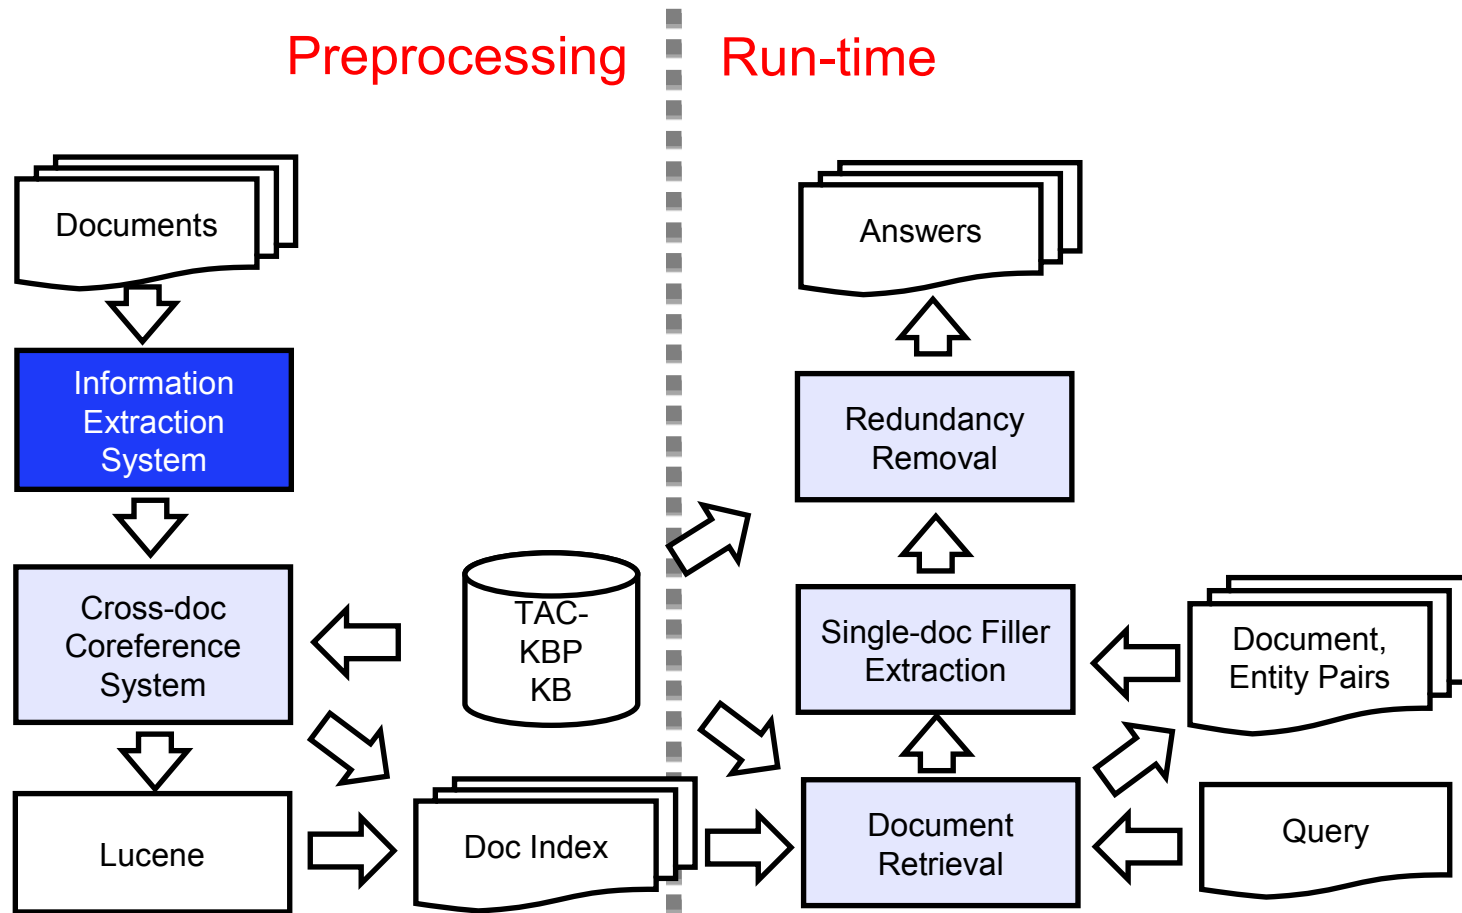
Overview

- **System descriptions**
- **Results**
- **Follow-up experiments**

System




No external knowledge source is used.



Information Extraction (IE)

Wu Shu-chen, the former first wife, visited her husband Chen Shui-bian in detention in the morning.

She was accompanied by their son Chen Chih-chung and Lawrence Gao, a Democratic Progressive Party lawmaker.



Information Extraction (IE)

PERSON

PEOPLE

ORGANIZATION

TIME

EVENT_MEETING

EVENT_CUSTOD
Y

Wu Shu-chen, the former first wife, visited her husband
Chen Shui-bian in detention in the morning.

She was accompanied by their son Chen Chih-chung
and Lawrence Gao, a Democratic Progressive Party
lawmaker.



Information Extraction (IE)

PERSON

PEOPLE

ORGANIZATION

TIME

EVENT_MEETING


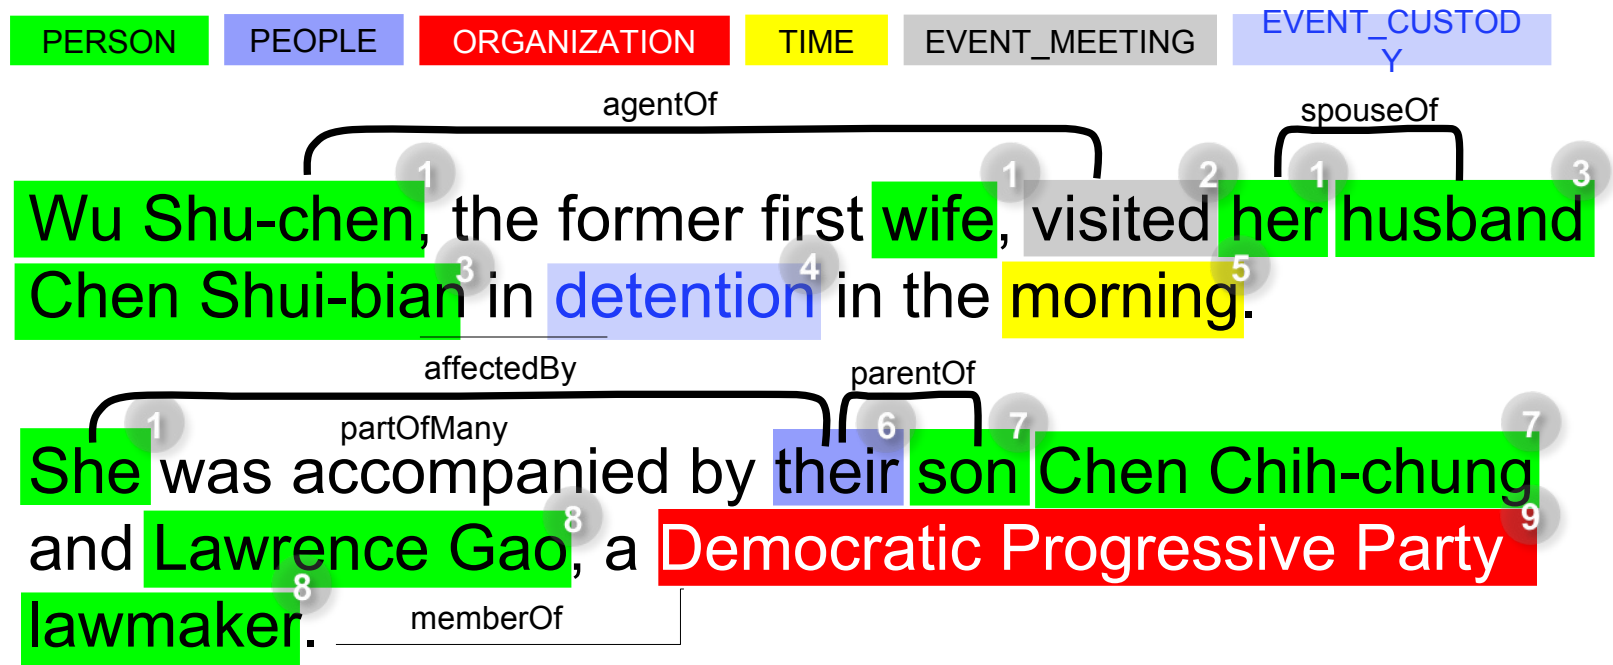
EVENT_CUSTODY

Wu Shu-chen, the former first wife, visited her husband Chen Shui-bian in detention in the morning.


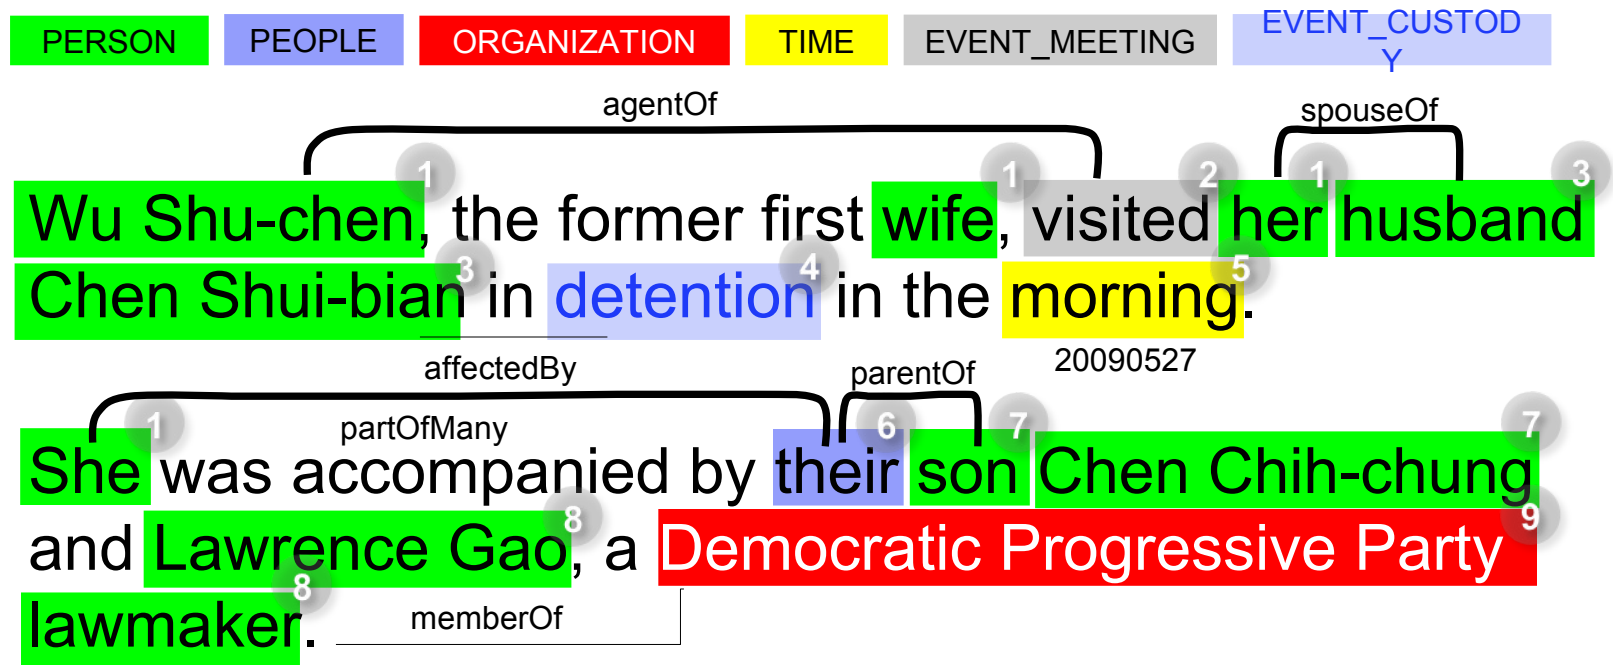
She was accompanied by their son Chen Chih-chung and Lawrence Gao, a Democratic Progressive Party lawmaker.



Information Extraction (IE)



Information Extraction (IE)



Information Extraction System

- **Trained on data annotated based on KLUE2 ontology**
 - 52 entity types: PERSON, ORG, EVENT_MEETING etc
 - 47 relation types: locatedAt, employeeOf, partOfMany etc
- **Targeted annotations for low-count relations (< 100 instances)**
 - 22 relation types
 - 846 additional documents (short self-sufficient paragraphs): 16.9k mentions and 8.3k relations
- **Total: 1367 documents**
 - 85.8k mentions and 33.4k relations

Improvements On Mention Detection/Coreference

■ **Mention Detection**

- Using revised annotation ontology KLUE2
- Improved from last year: $F = 78.57 \rightarrow 82.95$

■ **Coreference Resolution**

- Used hard constraints induced by parse-tree paths
- Improved from last year: $F = 68.48 \rightarrow 69.31$ (on system mentions)
- While reducing run-time by 50%

Improvements on Relation Detection

■ Sequential Decoding of Relations


- Modeling the dependency between relations
- Are both **locatedAt** relations valid?

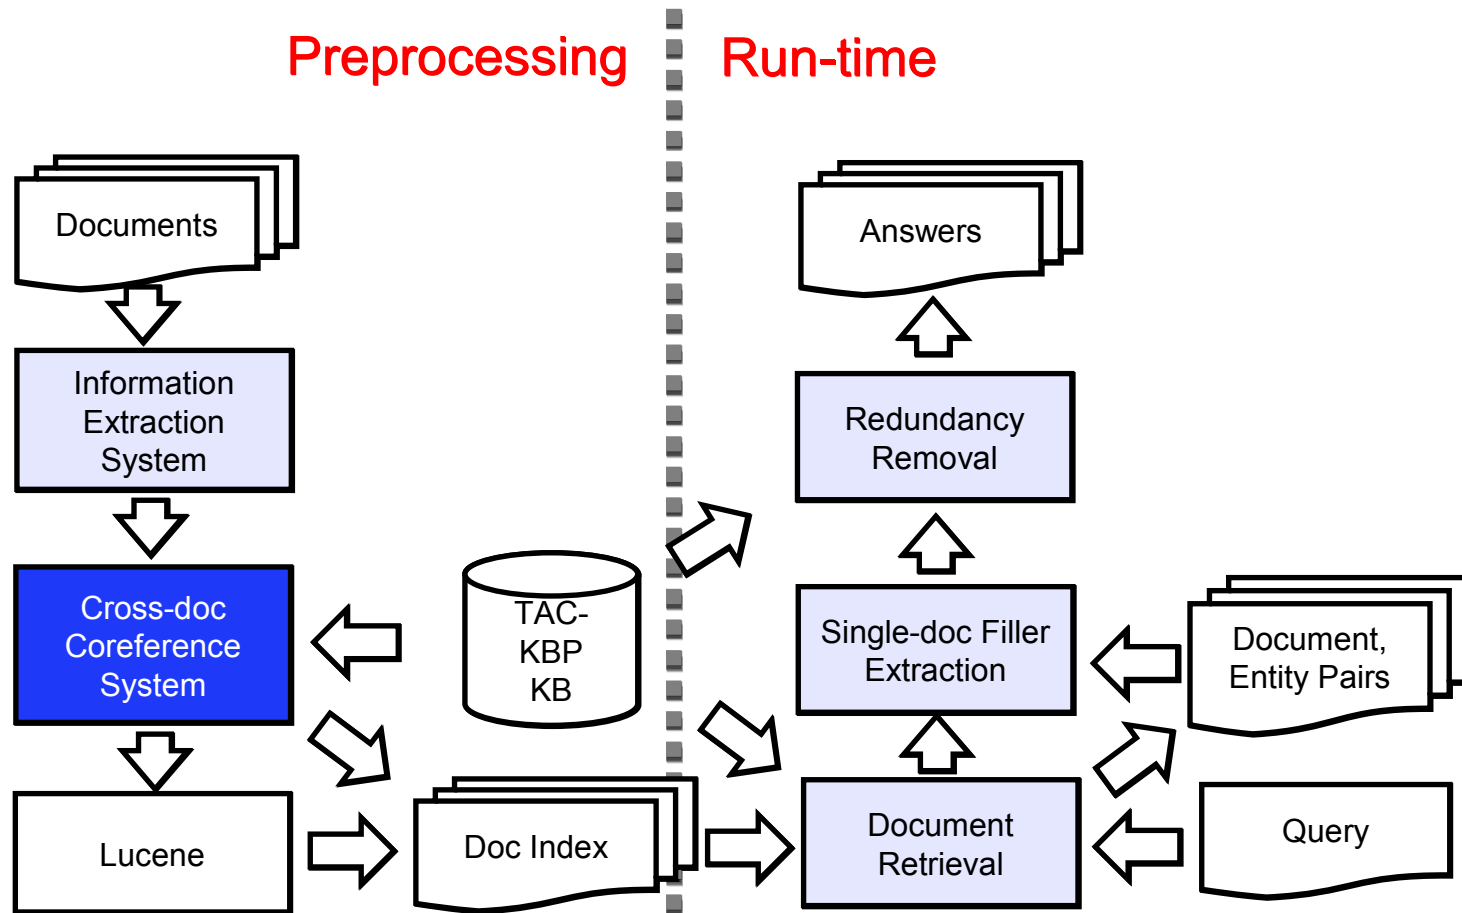
- A 23-year-old **man** will appear in **court** Thursday in connection with the failed bombings in **London**.

■ Decode using a stack decoder

- Order mention pairs within each sentence, from left to right
- For each mention pair, run the existence detector and the type classifier (both are MaxEnt-based)

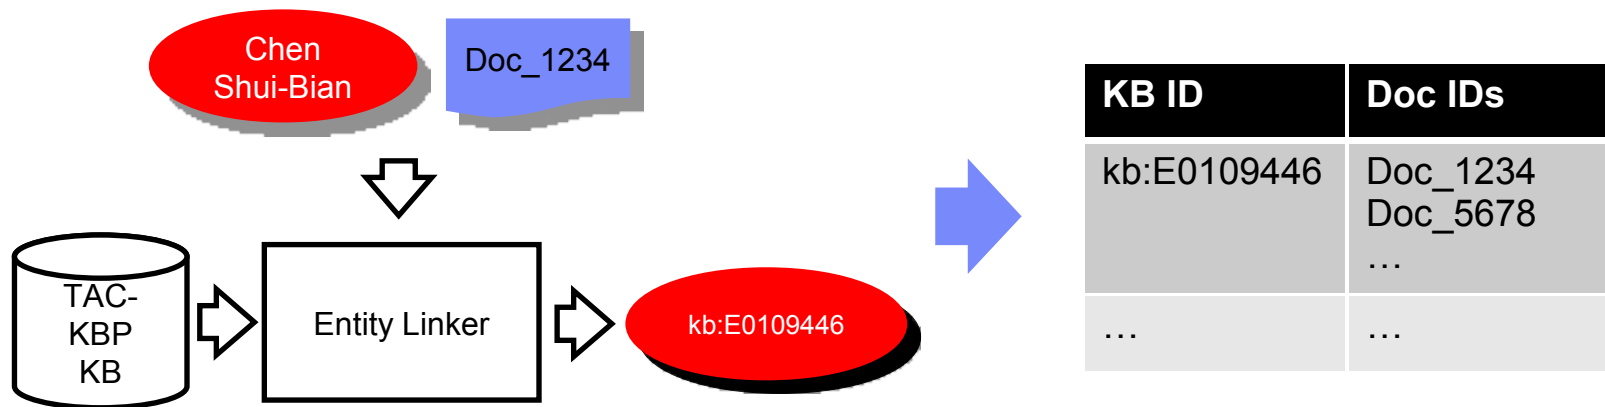
Relation Detection

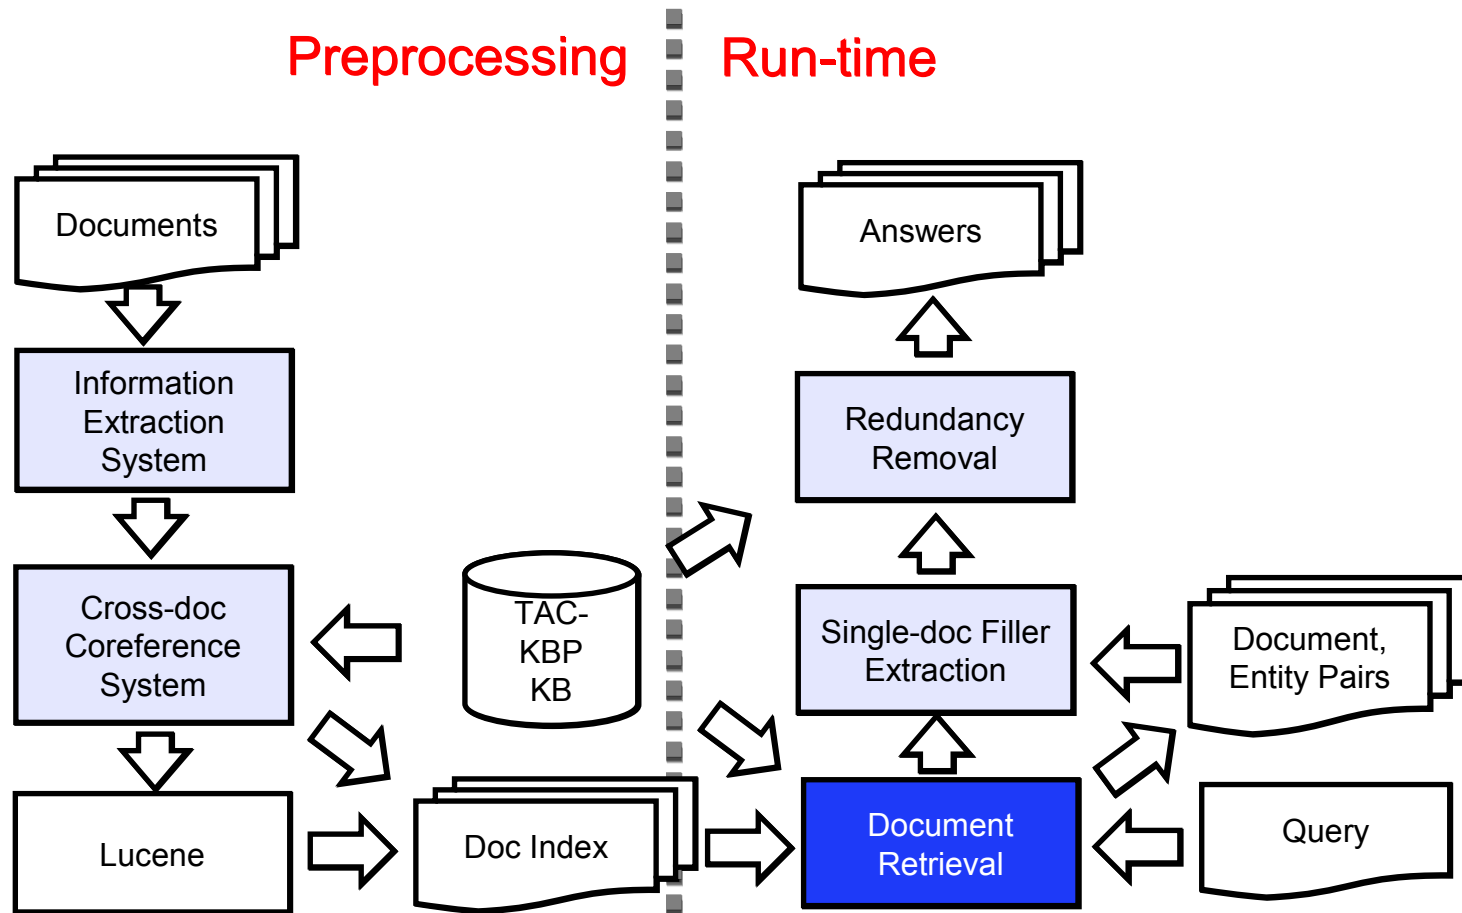




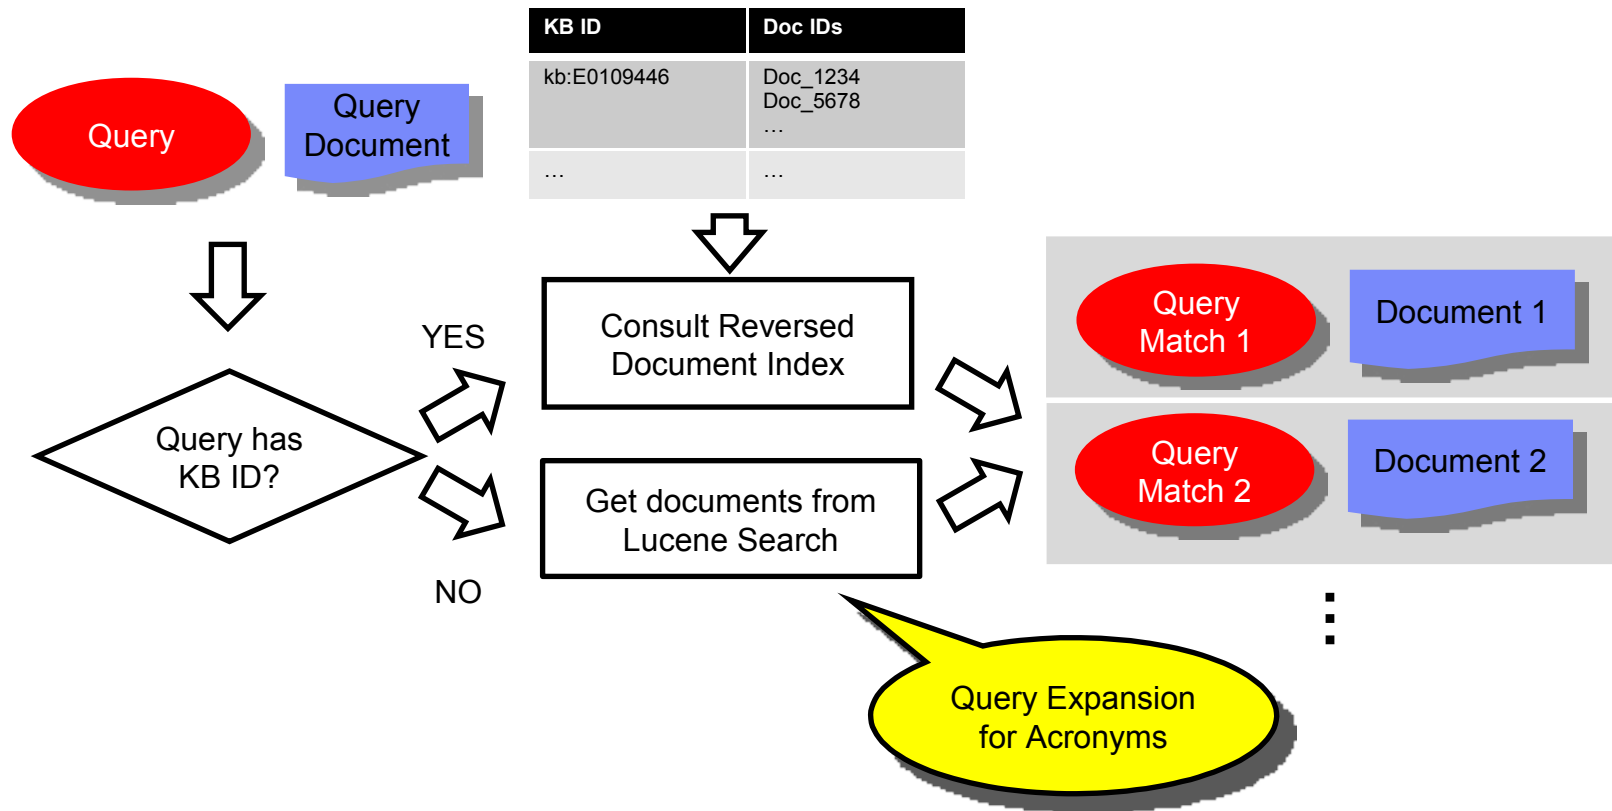
Cross-Document Coreference

- **Via Entity Linking** (TAC-KBP 2009)
 - (Entity, Document) → KB ID or null
 - Built a reverse document index
 - Keys are KB ID, values are documents containing the KB ID





Document Retrieval

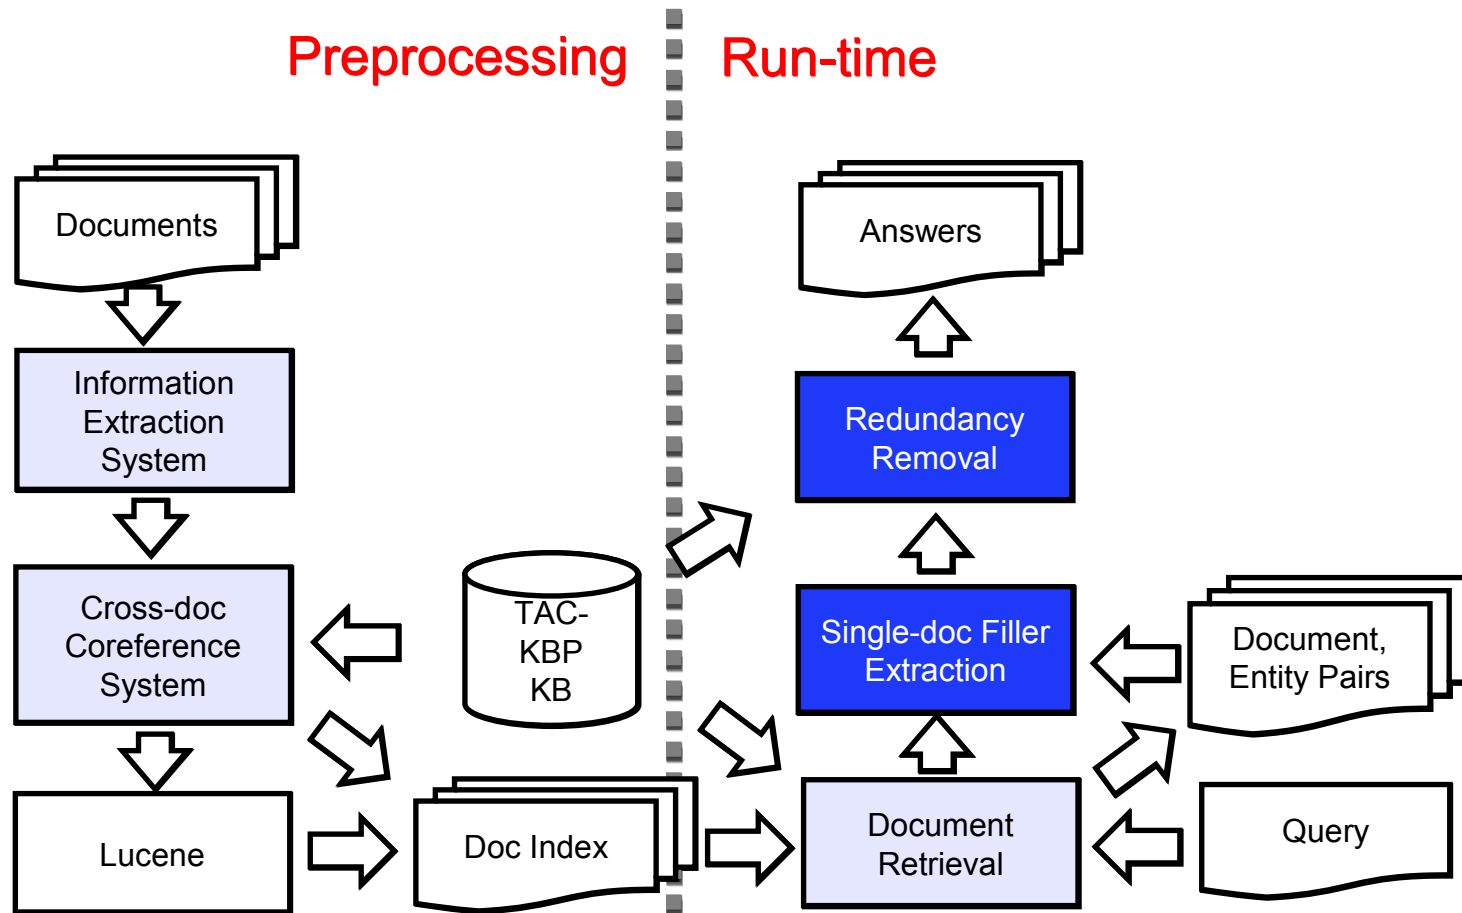


Documents Indexing with Lucene

- **Indexing only mention strings**
 - If we miss mentions, we miss documents
- **Improving mention detection**
 - Added query mentions to dictionaries of applicable types
 - “CC” is added to the ORGANIZATION dictionary
 - Nudge the system to treat them like the other dictionary entries
 - Improved document recall: 77.79 → 84.62 (LDC training queries)

Acronym Query: “CC”

- **Comedy Central? Circuit City?**
 - Query Document: *Following an investigation, the **Competition Commission** (CC) said it was seeking ...*
- **If a query is an acronym**
 - Find full names in the query document
 - Retrieve documents using both the acronym and the full names
- **Significant improvement: doc recall 0 → 100 while reducing # of docs retrieved from 1.3k to 180.**



Example: per:children

Query: Chen Shui-bian

Wu Shu-chen, the former first wife, visited her husband Chen Shui-bian in detention in the morning.

She was accompanied by their son Chen Chih-chung and Lawrence Gao, a Democratic Progressive Party lawmaker.

Example: per:children

Query: Chen Shui-bian

Wu Shu-chen, the former first wife, visited her husband Chen Shui-bian in detention in the morning.

The diagram illustrates a semantic relationship between the words 'her' and 'husband' in the sentence. A bracket labeled 'per:spouse' connects the two words. The word 'her' is annotated with a '1' and a '3' inside a grey circle, while 'husband' is annotated with a '3' inside a grey circle. The phrase 'Chen Shui-bian' is also annotated with a '3' inside a grey circle.

She was accompanied by their son Chen Chih-chung and Lawrence Gao, a Democratic Progressive Party lawmaker.

PERSON

PEOPLE

Example: per:children

Query: Chen Shui-bian

Wu Shu-chen, the former first wife, visited **her** **husband** **Chen Shui-bian** in detention in the morning.

She was accompanied by **their** son Chen Chih-chung and Lawrence Gao, a Democratic Progressive Party lawmaker.

PERSON

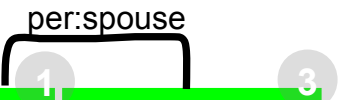
PEOPLE

Example: per:children

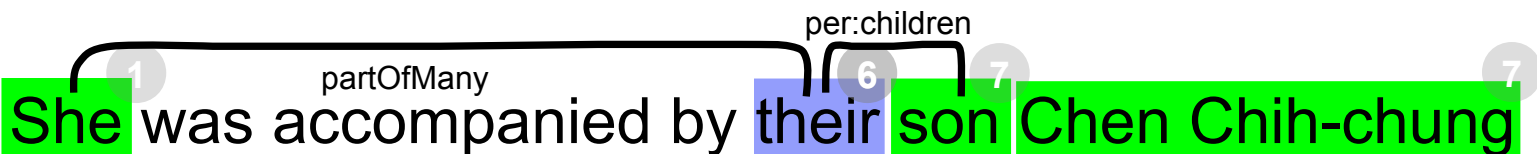
Query: Chen Shui-bian

Per:children: Chen Chih-chung

Wu Shu-chen, the former first wife, visited **her** **husband** **Chen Shui-bian** in detention in the morning.



She was accompanied by **their** **son** **Chen Chih-chung** and Lawrence Gao, a Democratic Progressive Party lawmaker.



PERSON

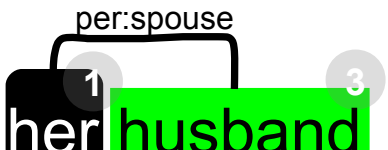
PEOPLE

Example: per:children

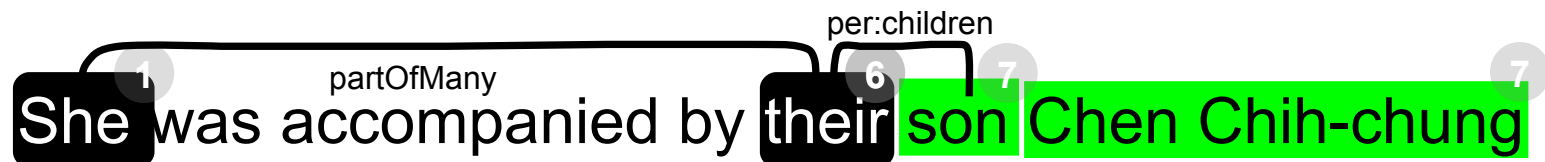
Query: Chen Shui-bian

Per:children: Chen Chuh-chung

Wu Shu-chen, the former first wife, visited **her** **husband** **Chen Shui-bian** in detention in the morning.



She was accompanied by **their** **son** **Chen Chih-chung** and Lawrence Gao, a Democratic Progressive Party lawmaker.

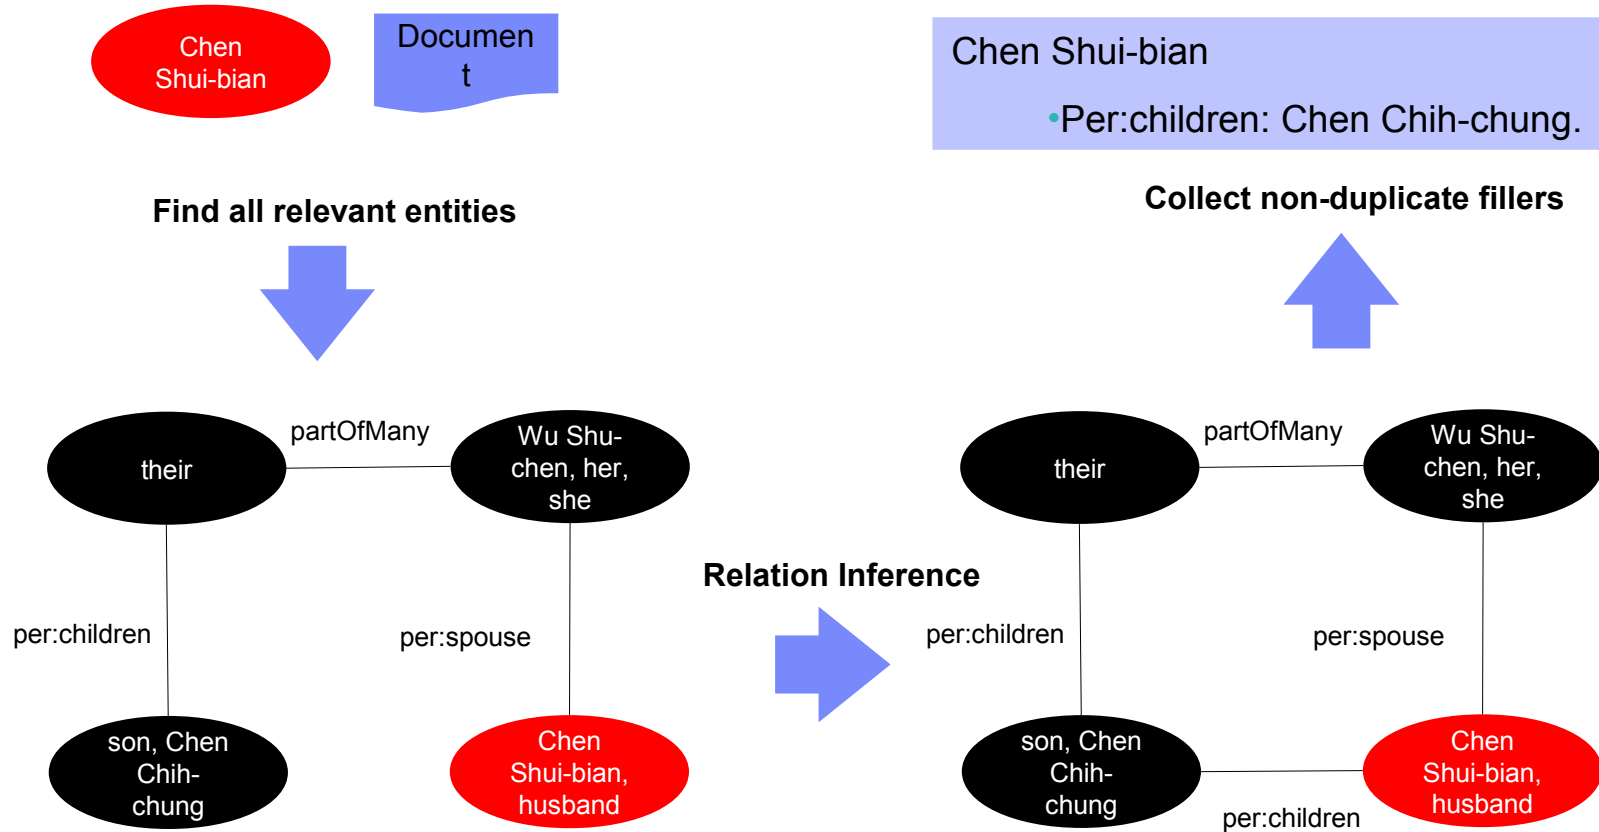


PERSON

PEOPLE

“third party” entities

Single-Document Slot Filler Extraction



Relation Inference

- **Level 0: no inference (just use the original relations)**
- **Level 1: use basic relation properties**
 - Equivalence, transitivity, symmetry.
- **Level 2: use simple implications**
 - *Palmisano* **managerOf** *IBM* implies *Palmisano* **employeeOf** *IBM*.
- **Level 3: relation chaining**
 - If *Ben* **colleague** *Vittorio* and *Vittorio* **employeeOf** *IBM*
then *Ben* **employeeOf** *IBM*.
- **Level 4: recursive reasoning (using extracted slots in inference)**
 - If *Chen Shui-bian* **per:spouse** *Wu* and *Wu* **per:children** *Chen Chih-chung*, then *Chen Shui-bian* **per:children** *Chen Chih-chung*.

Example Rules

per:date_of_birth(X,Y) :- bornOn(X,Y).
 per:age(X,Y) :- ageOf(X,Y).
 per:employee_of(X,Y) :- employedBy(X,Y).

120 Rules


per:religion(X,Y) :- partOfMany(X,Y), religious(Y).
 per:religion(X,Y) :- located(X,Z), religiousFacility(Z,Y).

per:origin(X,Y) :- isOrigin(Y), coref(X,Y).
 per:origin(X,Y) :- partOfMany(X,Y), isGPE(Y).
 per:origin(X,Y) :- per:sibling(X,Z), per:origin(Z,Y).
 per:origin(X,Y) :- partOfMany(X,Z), per:origin(Z,Y).

per:siblings(X,Y) :- isSibling(Y), relativeOf(X, Y).

per:siblings(X,Y) :- per:parents(Z,X), per:parents(Z,Y), X!=Y.

per:siblings(X,Y) :- partOfMany(X,Z), per:siblings(Z,Y), X!=Y.




Redundancy Removal

- **Accumulate instance counts for fillers.**
- **Group fillers into equivalence classes**
 - Fillers linked to the same KB entity are grouped.
 - Group fillers based on string similarity.
 - Heuristics for person/organization names.
- **Pick the highest n classes based on counts: representatives (longest) are answers.**

Results & Conclusions

Internal Results on LDC Training Queries




Official Results

- **IBM1: max 5k docs for extraction, and 2 slot types filtered**
(per:employee_of and per:charges)
- **IBM2: Same as IBM1 but no slot type was filtered**
- **IBM3: Same as IBM1 but with max 800 docs for extraction**

System	Precision	Recall	F-measure
IBM1	28.0	27.0	27.5
IBM2	25.3	29.0	27.0
IBM3	31.0	25.9	28.2
Human	70.1	54.1	61.1
TopSystem	66.8	64.8	65.8
Top2System	66.5	18.7	29.2

Effect of Inference on Performance



Level 0: no inference
Level 1: basic relation properties
Level 2: simple relation implications

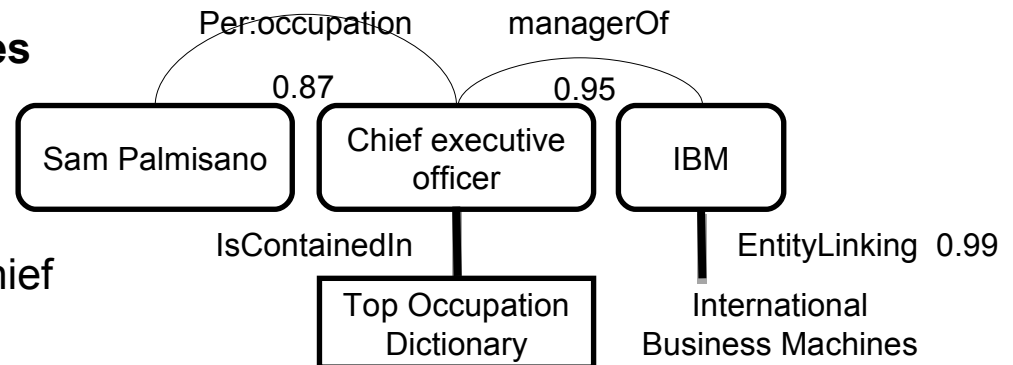
Level 3: relation chaining
Level 4: recursive reasoning

Post-filtering


- **Reject a filler based on its inference traces**
 - Potential gain in precision
- **Trained a MaxEnt model to reject a filler**
 - Training data: system output without redundancy checks
 - Features: inference traces with no lexical info
 - 3.7k/ 4.5k positive/negative examples

Org:top_members/employees

“Sam Palmisano, the current chief executive officer of IBM, ...”



Post-filtering (10-fold cross-validation)



Less conservative

	Accuracy	True Positive %	False Negative %	False Positive %	True Negative %
T = 1	33.4	33.4	0	66.7	0
T = 0.8	56.7	31.3	2.0	41.3	25.4
T = 0.5	73.3	19.4	12.7	14.0	53.9

Reject a filler only if classifier predicts **WRONG** with confidence $> T$

Conclusions

- **Demonstrated an effective combination of statistical IE and rule-based reasoning**
- **Observed significant benefits from recursive reasoning**
- **Established a direction towards a fully statistical system**